Average performance of Orthogonal Matching Pursuit (OMP) for sparse approximation

Karin Schnass   KARIN.SCHNASS@UIBK.AC.AT

Department of Mathematics
University of Innsbruck
Technikerstraße 13
6020 Innsbruck, Austria

Abstract

We present a theoretical analysis of the average performance of OMP for sparse approximation. For signals that are generated from a dictionary with K atoms and coherence μ and coefficients corresponding to a geometric sequence with parameter $\alpha < 1$, we show that OMP is successful with high probability as long as the sparsity level S scales as $S\mu^2 \log K \lesssim 1 - \alpha$. This improves by an order of magnitude over worst case results and shows that OMP and its famous competitor Basis Pursuit outperform each other depending on the setting.

Keywords: sparse approximation; Orthogonal Matching Pursuit; average case analysis; decaying coefficients;

1. Introduction

In sparse approximation the goal is to approximate a given signal $y \in \mathbb{R}^d$ by a linear combination of a small number $S \ll d$ of elements $\varphi_i \in \mathbb{R}^d$, called atoms, out of a given larger set, such as basis or a frame, called the dictionary. Storing the normalised atoms as columns in the dictionary matrix $\Phi = (\varphi_1 \ldots, \varphi_K)$, and denoting the restriction to the columns indexed by a set I by Φ_I, we can write informally,

$$\text{find} \quad y \approx \sum\nolimits_{k \in I} \varphi_k x_k = \Phi_I x_I \quad \text{s.t.} \quad |I| = S \ll d \tag{1}$$

Finding the smallest error for a given sparsity level S and the corresponding support set I, which determines x_I via $x_I = \Phi_I^\dagger y$, where $\Phi_I^\dagger$ is the Moore-Penrose pseudo inverse, becomes an NP-hard problem in general unless the dictionary is an orthonormal system. In this case thresholding, meaning choosing as I the indices of the atoms having the S-largest inner products with the signal in magnitude, will succeed. For all other cases, one had to find algorithms which are more efficient, if less optimal than an exhaustive search through all possible supports sets I with subsequent projection $P(\Phi_I)y := \Phi_I \Phi_I^\dagger y$. The two most investigated directions are greedy methods and convex relaxation techniques - the two golden classics being Orthogonal Matching Pursuit (OMP), [13], and Basis Pursuit (BP), [3], respectively.
OMP finds the support iteratively, adding the index of the atom which has the largest absolute inner product with the residual and updating the residual. So initialising $r_0 = y$,

$J_0 = \emptyset$, it

$$\text{finds} \quad j = \text{argmax}_k \, |\langle r_i, \varphi_k \rangle| \quad \text{and}$$
$$\text{updates} \quad J_{i+1} = J_i \cup \{j\} \quad \text{resp.} \quad r_{i+1} = y - P(\Phi_{J_{i+1}})y,$$

until a stopping criterion is met, such as reaching the desired number of iterations or the size of the residual/largest inner product being sufficiently small.

The Basis Pursuit principle, on the other hand, prescribes finding the minimiser of the convex programme

$$\hat{x} = \text{argmin}_{x:y=\Phi x} \|x\|_1, \tag{2}$$

and choosing I as the index set of the S-largest entries of $\hat{x}$ in magnitude.

The interesting question concerning both schemes is when they are successful. So assuming that the signal y is known to be S-sparse, meaning $y = \Phi_I x_I$ with $|I| = S$, when can they recover the support I. It was first studied in [15, 7] and for dictionaries with coherence $\mu := \max_{j \neq k} |\langle \varphi_j, \varphi_k \rangle|$ a sufficient condition for both schemes to succeed is that $2S\mu < 1$, which is relaxed in comparison to the sufficient condition for thresholding $2S\mu < \frac{\min_{k \in I} |x_k|}{\max_{k \in I} |x_k|}$, but still quite restrictive, especially considering the much better performance in practice. This led to the investigation of the average performance when modelling the signals as

$$y = \sum_k \sigma_k c_k \varphi_{p(k)}, \tag{3}$$

where $(\sigma_k)_k$ is a Rademacher sequence, the coefficient sequence c is non-increasing, $c_k \geq c_{k+1} \geq 0$, and $c_k = 0$ for $k > S$ and p is some permutation such that the support $I = \{p(1), \ldots, p(S)\})$ satisfies $\delta_I := \|\Phi_I^\star \Phi_I - \mathbb{I}_d\|_{2,2} \leq \frac{1}{2}$, where for a matrix A the transpose is denoted by $A^\star$.

It was shown that BP recovers the true support except with probability $2K^{1-2m}$ as long as $16\mu^2 S \cdot m \log K \leq 1$,[16][1], and that thresholding succeeds except with probability $2K^{1-2m}$ as long as $32\mu^2 S \cdot m \log K \leq \frac{\min_{k \in I} |x_k|}{\max_{k \in I} |x_k|}$, [14][2]. The fact that for OMP a similar result could only be found in a multi-signal scenario, [9], started to give OMP the reputation of being weaker than BP.

This was further increased by the advent of Compressed Sensing (CS), [4], which can be seen as sparse approximation with design freedom for the dictionary. While for BP-type schemes in combination with randomly chosen dictionaries strong results appeared very early, [2, 1], comparable results for OMP and its variants took longer to develop and are weaker in general, [8, 12]. Still, thanks to its computational advantages and flexibility, e.g. concerning the stopping criteria, OMP remained popular in signal processing - the only difference being that users had a defensive statement a la 'of course BP will perform even better' ready at all times.

Contribution: Here we will provide the long missing analysis of the average performance of OMP and show that on average neither BP nor OMP are stronger, but confirm folklore

---

1. The theorem actually considers Steinhaus instead of Rademacher sequences. However, the proof for Rademacher sequences is exactly the same; simply use Hoeffding's inequality instead of the complex Bernstein inequality. Also beware the buggy support condition in model M1.
2. The improved constant presented here is due to the fact that also for Rademacher sequences $c_0 = \frac{1}{2}$.

wisdom, that OMP works better for signals with decaying coefficients while BP is better for equally sized coefficients. The idea that the performance of OMP improves for decaying coefficients has already been used in [10] and the simplified result states that if the sorted absolute coefficients form a geometric sequence with decay $\alpha \leq \frac{1}{2}$, then OMP is guaranteed to succeed for all sparsity levels S with $S\mu < 1$. Replacing certainty with high probability we will relax this bound by an order of magnitude to $S\mu^2 \sqrt{\log K} \lesssim 1 - \alpha$, for $\alpha < 1$.

Organisation: We will first address full support recovery in the noiseless case and then extend this result to partial support recovery in the noisy case. In Section 4 we will conduct two experiments showing that our theoretical results accurately predict the average performance, before finally discussing our results and future work.

2. Noiseless Case

We start with the simple case of signals following the model in (3). Note that from [16] we know that for a randomly chosen subset I (permutation p) the condition $\delta_I \leq \frac{1}{2}$ is satisfied with high probability as long as $\mu^2 S \log K \lesssim 1$.

Theorem 1 *Assume that the signals follow the model in (3) and that for $i \leq S$ the coefficients satisfy $c_{i+t}/c_i \leq 1 - \frac{\lambda}{S}$ for $t, \lambda > 0$. Then, except with probability $2SK^{1-2m}$, OMP will recover the full support as long as*

$$\left(t \left\lceil \frac{S}{\lambda} \right\rceil + \sqrt{\frac{mtS \log K}{\lambda}} + 1 \right) S\mu^2 \leq \frac{1}{13}. \tag{4}$$

Before presenting the proof, we want to provide some background information on the ideas used for proving success of OMP and the difficulties associated with an average case analysis. A necessary and sufficient condition for a step of OMP to succeed is that for the current (correct) sub-support $J \subset I$ we have

$$\max_{i \in I} |\langle \varphi_i, r_J \rangle| > \max_{k \notin I} |\langle \varphi_k, r_J \rangle|. \tag{5}$$

Thus a sufficient condition for OMP to fully recover the support is, that for all possible sub-supports J the missing atom which has the largest coefficient φ_{i_J} satisfies

$$|\langle \varphi_{i_J}, r_J \rangle| > \max_{k \notin I} |\langle \varphi_k, r_J \rangle|. \tag{6}$$

If the coefficients have random signs then for all k, J the inner products should concentrate around their expectation,

$$|\langle \varphi_i, r_J \rangle|^2 \rightsquigarrow \sum_{k \in I/J} |\langle \varphi_i, [\mathbb{I}_d - P(\Phi_J)] \varphi_k \rangle|^2 c_k^2 \rightsquigarrow c_i^2 \pm \mu^2 \|c_{I/J}\|_2^2, \tag{7}$$

so a condition of the form $S\mu^2 \lesssim 1$ should ensure success with high probability. The problem is that there are 2^S sub-supports J for which we need to have this concentration. So taking a union bound for the probability of not enough concentration over all sub-supports, we get back the worst case condition but with a non-zero failure probability. The immediate conclusion is that in order to get a useful average case result, we have to reduce the

number of intermediate supports that we need to control. For equally sized coefficients this is impossible, since the random signs determine the order of the absolute inner products. However, if the coefficients exhibit some decay, there is a natural order and it is more likely that atoms with large coefficients are picked first. For instance, with sufficient decay it might happen that the atom with the second largest coefficient is picked before that with the largest, but very unlikely that the atom with the smallest coefficient is picked first. The idea of the proof is that OMP will only pick 'sensible' sub-supports, so we only need to ensure concentration for a much smaller number of them. The amount of concentration needed can then be further reduced by pooling 'sensible' supports of the same type and combining probabilistic and deterministic bounds.

Proof We use the following short hands $Q(\Phi_J) = \mathbb{I}_d - P(\Phi_J)$ as well as $r_J = Q(\Phi_J)y$ for the residual based on an index set J and x for the signed coefficients, $x_k := c_k \sigma_k$. In order to better understand the various bounds of terms involving Φ_J, we recommend a quick familiarisation with Lemma 6.2, [9]. Further we will assume w.l.o.g., that is, by reordering the dictionary matrix, that $I = \{1, \ldots, S\} =: \mathbb{S}$. We now define the following disjoint sets for $i \leq S$ and a parameter $T > 0$, which we set to the optimal value later,

$$A_i = \{1 \ldots i - 1\} \quad \text{with} \quad A_1 = \emptyset, \tag{8}$$

$$M_i = \{i + 1, \ldots, i + T - 1\} \cap \mathbb{S}, \tag{9}$$

$$Z_i = \{i + T, \ldots, K\} \cup \mathbb{S}^c. \tag{10}$$

We call a sub-support $J \subseteq \mathbb{S}$ admissible if there exist $i \leq S$, corresponding to the index of the missing atom with largest coefficient in magnitude, and $B \subseteq M_i$ such that $J = A_i \cup B$. We write $B^c := M_i/B$. Note that an admissible sub-support $J = A_i \cup B$ contains the indices of the atoms with the first $i - 1$ largest coefficients in magnitude, does not contain the index i, may contain indices in the support corresponding to atoms with coefficients large enough that they are likely to be picked before i and is not allowed to contain any indices corresponding to atoms with too small coefficients or outside the support.

A sufficient condition for OMP to succeed is that it only picks admissible sub-supports. Assuming J is admissible, OMP picks another admissible support if (suff. cond.)

$$|\langle \varphi_i, r_J \rangle| > \max_{k \in Z_i} |\langle \varphi_k, r_J \rangle|, \tag{11}$$

which ensures that the either i or some $k \in B^c$ is chosen. Since J is admissible the residual has the form

$$r_J = Q(\Phi_J)y = Q(\Phi_J)(x_i \varphi_i + \Phi_{B^c} x_{B^c} + \Phi_{Z_i} x_{Z_i}), \tag{12}$$

and we have for i, the index of the largest missing coefficient,

$$\langle \varphi_i, r_J \rangle = x_i \|Q(\Phi_J)\varphi_i\|_2^2 + \langle \varphi_i, Q(\Phi_J)\Phi_{B^c} x_{B^c} \rangle + \langle \varphi_i, Q(\Phi_J)\Phi_{Z_i} x_{Z_i} \rangle$$
$$= x_i \|Q(\Phi_J)\varphi_i\|_2^2 + \langle \varphi_i, Q(\Phi_J)\Phi_{B^c} x_{B^c} \rangle + \langle \varphi_i, \Phi_{Z_i} x_{Z_i} \rangle - \langle \varphi_i, P(\Phi_J)\Phi_{Z_i} x_{Z_i} \rangle. \tag{13}$$

For $k \in Z_i$ we define $Z_i^k := Z_i/\{k\}$ and can rewrite (12) as

$$r_J = Q(\Phi_J)y = Q(\Phi_J)(x_i \varphi_i + \Phi_{B^c} x_{B^c} + x_k \varphi_k + \Phi_{Z_i^k} x_{Z_i^k}),$$

4

which leads to

$$\langle \varphi_k, r_J \rangle = x_i \langle \varphi_k, Q(\Phi_J)\varphi_i \rangle + \langle \varphi_k, Q(\Phi_J)\Phi_{B^c} x_{B^c} \rangle$$
$$+ x_k + \langle \varphi_k, \Phi_{Z_i^k} x_{Z_i^k} \rangle - \langle \varphi_k, P(\Phi_J)\Phi_{Z_i} x_{Z_i} \rangle. \tag{14}$$

We first bound the terms in (13/14) involving $Q(\Phi_J)$. So for all $k \in Z_i \cup \{i\}$ we have

$$\begin{aligned}
\|Q(\Phi_J)\varphi_k\|_2^2 &= 1 - \|P(\Phi_J)\varphi_k\|_2^2 \\
&= 1 - \langle \Phi_J^\star \varphi_k, (\Phi_J^\star \Phi_J)^{-1} \Phi_J^\star \varphi_k \rangle \\
&\geq 1 - \|\Phi_J^\star \varphi_k\|_2 \cdot \|(\Phi_J^\star \Phi_J)^{-1}\|_{2,2} \cdot \|\Phi_J^\star \varphi_k\|_2 \geq 1 - \frac{|J|\mu^2}{1 - \delta_J},
\end{aligned} \tag{15}$$

where we have used the bound $\|(\Phi_J^\star \Phi_J)^{-1}\|_{2,2} \leq (1 - \delta_J)^{-1}$ from Lemma 6.2 in [9]. Similarly for all $k \in Z_i$ we have

$$\begin{aligned}
|\langle \varphi_k, Q(\Phi_J)\varphi_i \rangle| &= |\langle \varphi_k, \varphi_i \rangle - \langle \Phi_J^\star \varphi_k, (\Phi_J^\star \Phi_J)^{-1} \Phi_J^\star \varphi_i \rangle| \\
&\leq \mu + \|\Phi_J^\star \varphi_k\|_2 \cdot \|(\Phi_J^\star \Phi_J)^{-1}\|_{2,2} \cdot \|\Phi_J^\star \varphi_i\|_2 \leq \mu + \frac{|J|\mu^2}{1 - \delta_J}.
\end{aligned} \tag{16}$$

and again for all $k \in Z_i \cup \{i\}$,

$$\begin{aligned}
|\langle \varphi_k, Q(\Phi_J)\Phi_{B^c} x_{B^c} \rangle| &\leq |\langle \varphi_k, \Phi_{B^c} x_{B^c} \rangle| + |\langle \varphi_k, P(\Phi_J)\Phi_{B^c} x_{B^c} \rangle| \\
&\leq \|\Phi_{B^c}^\star \varphi_k\|_\infty \cdot \|x_{B^c}\|_1 + \|\Phi_{B^c}^\star P(\Phi_J)\varphi_k\|_\infty \cdot \|x_{B^c}\|_1 \\
&\leq \|x_{B^c}\|_1 \cdot \left(\mu + \max_{j \in B^c} |\langle \varphi_j, P(\Phi_J)\varphi_k \rangle| \right) \\
&\leq \|x_{B^c}\|_1 \left(\mu + \frac{|J|\mu^2}{1 - \delta_J} \right).
\end{aligned} \tag{17}$$

The terms involving $P(\Phi_J)$ can be bounded for all $k \in Z_i \cup \{i\}$ as

$$\begin{aligned}
|\langle \varphi_k, P(\Phi_J)\Phi_{Z_i} x_{Z_i} \rangle| &= \langle \Phi_J^\star \varphi_k, (\Phi_J^\star \Phi_J)^{-1} \Phi_J^\star \Phi_{Z_i} x_{Z_i} \rangle \\
&\leq \|\Phi_J^\star \varphi_k\|_2 \cdot \|(\Phi_J^\star \Phi_J)^{-1}\|_{2,2} \cdot \|\Phi_J^\star \Phi_{Z_i} x_{Z_i}\|_2 \\
&\leq \frac{\sqrt{|J|}\mu}{1 - \delta_J} \cdot \|\Phi_J^\star \Phi_{Z_i} x_{Z_i}\|_2 \\
&\leq \frac{|J|\mu}{1 - \delta_J} \cdot \max_{j \in J} |\langle \varphi_j, \Phi_{Z_i} x_{Z_i} \rangle| \leq \frac{|J|\mu}{1 - \delta_J} \cdot \max_{j \notin Z_i} |\langle \varphi_j, \Phi_{Z_i} x_{Z_i} \rangle|.
\end{aligned}$$

From the bound above we can see that the only terms we still need to control are $|\langle \varphi_j, \Phi_{Z_i} x_{Z_i} \rangle|$ for $j \notin Z_i$ and $|\langle \varphi_k, \Phi_{Z_i^k} x_{Z_i^k} \rangle|$ for $k \in Z_i$. Note that for $j \notin Z_i$ we have $Z_i^j = Z_i$, meaning via Hoeffding's inequality we can estimate compactly for all k,

$$\begin{aligned}
\mathbb{P}\left(|\langle \varphi_k, \Phi_{Z_i^k} x_{Z_i^k} \rangle| \geq \theta_k \right) &= \mathbb{P}\left(\left| \sum_{j \in Z_i^k} \langle \varphi_k, \varphi_j \rangle c_j \sigma_j \right| > \theta_k \right) \\
&\leq 2 \exp\left(\frac{-\theta_k^2}{2 \sum_{j \in Z_i^k} |\langle \varphi_k, \varphi_j \rangle|^2 c_j^2} \right) \leq 2 \exp\left(\frac{-\theta_k^2}{2\mu^2 \|c_{Z_i^k}\|_2^2} \right).
\end{aligned}$$

Setting $\theta_k = 2\theta\mu\|c_{Z_i}\|_2$ and using a union bound, we get

$$\mathbb{P}\left(\exists k : |\langle\varphi_k, \Phi_{Z_i^k} x_{Z_i^k}\rangle| \geq 2\theta\mu\|c_{Z_i}\|_2\right) \leq 2Ke^{-2\theta^2}. \tag{18}$$

Combining all our estimates we get that except with probability $2Ke^{-2\theta^2}$ for all $k \in Z_i$

$$|\langle\varphi_i, r_J\rangle| - |\langle\varphi_k, r_J\rangle| \geq c_i\left(1 - \mu - \frac{2|J|\mu^2}{1-\delta_J}\right) - c_k - 2\|c_{B^c}\|_1\left(\mu + \frac{|J|\mu^2}{1-\delta_J}\right)$$
$$- 4\theta\|c_{Z_i}\|_2\left(\mu + \frac{|J|\mu^2}{1-\delta_J}\right).$$

We now determine the sets M_i and Z_i by choosing $T := t \cdot \lceil\frac{S}{\lambda}\rceil$ to get the following bounds. For all $k \in Z_i$ we have

$$c_k \leq c_i\left(1 - \frac{\lambda}{S}\right)^{T/t} \leq c_i\left(1 - \frac{\lambda}{S}\right)^{S/\lambda} \leq c_ie^{-1}. \tag{19}$$

We also have $\|c_{B^c}\|_1 \leq \|c_{M_i}\|_1 \leq c_it\lceil\frac{S}{\lambda}\rceil$. To bound $\|c_{Z_i}\|_2$ we note that at worst c_{Z_i} consists of t interleaved geometric sequences with initial value c_ie^{-1} and decay factor $\alpha = 1 - \lambda/S$, so we get

$$\|c_{Z_i}\|_2 \leq c_ie^{-1}\left(\frac{t}{1-(1-\lambda/S)^2}\right)^{1/2} \leq c_ie^{-1}\left(\frac{tS}{\lambda}\right)^{1/2}.$$

Using $\delta_J \leq \delta_I \leq \frac{1}{2}$ and setting $\theta := \sqrt{m\log K}$, we get that except with probability $2K^{1-2m}$ for all $k \in Z_i$,

$$c_i^{-1}\left(|\langle\varphi_i, r_J\rangle| - |\langle\varphi_k, r_J\rangle|\right) \geq 1 - \mu - 4|J|\mu^2 - e^{-1} - 2t\left\lceil\frac{S}{\lambda}\right\rceil\left(\mu + 2|J|\mu^2\right)$$
$$- 4e^{-1}\sqrt{\frac{mtS\log K}{\lambda}}\left(\mu + 2|J|\mu^2\right) \tag{20}$$
$$> 0.63 - 2\left(t\left\lceil\frac{S}{\lambda}\right\rceil + \sqrt{\frac{mtS\log K}{\lambda}} + 1\right)\left(\mu + 2S\mu^2\right).$$

In case $2S\mu \leq 1$ the deterministic analysis holds and the theorem is trivially true. If conversely $2S\mu \geq 1$, then $\mu + 2S\mu^2 \leq 4S\mu^2$ and so (4) implies that the last expression above is larger than zero, which further implies that OMP will pick another admissible sub-support. Taking a union bound over all possible sets A_i we get that OMP will succeed with probability at least $1 - 2SK^{1-2m}$ as long as (4) holds. ∎

First note that the theorem only improves over the worst case analysis when $\lambda > t$. However, if conversely $\lambda \leq t$, then the ratio between largest and smallest coefficient is of the order $\approx e^{-1}$, so thresholding should still have a good success probability.

To get a better feeling for the quality of the theorem, we next specialise it to the case $t = 1$, where the coefficients form a sub-geometric sequence with parameter α, meaning $c_{i+1}/c_i \leq \alpha < 1$. In this case the theorem essentially says that OMP will recover the support except

with probability $2SK^{1-2m}$ as long as $S\mu^2 \lesssim 1-\alpha$ and $S\mu^2\sqrt{m\log K} \lesssim \sqrt{1-\alpha}$. Comparing this to the condition for BP, $S\mu^2 m\log K \lesssim 1$, for failure probability $2K^{1-2m}$, we see that OMP has the advantage that the admissible sparsity level has a milder dependence on the dictionary size and success probability while BP has the advantage of being independent of the coefficient decay. This means that each algorithm can outperform the other depending on the setting. Before confirming this in the numerical simulation in Section 4 we first have a look at the performance of OMP in a noisy setting.

3. Noisy Case

We next study partial support recovery, when the sparse signals are contaminated with noise and are modelled as

$$\tilde{y} = y + \eta = \sum_k \sigma_k c_k \varphi_{p(k)} + \eta, \tag{21}$$

with y as in the previous section and η a sub-Gaussian noise vector with parameter ρ. This means that $\mathbb{E}(\eta) = 0$ and that for all unit vectors v and $\theta > 0$ the marginals $\langle v, \eta \rangle$ satisfy $\mathbb{E}(e^{\theta\langle v,\eta \rangle}) \le e^{\theta^2\rho^2/2}$.

For Gaussian noise the parameter ρ corresponds to the standard deviation and so for normalised coefficient sequences $\|c\|_2 = 1$ the signal to noise ratio (SNR) is $\frac{1}{d\rho^2}$. Similar bounds also hold in the general case, [11]. In the noisy setting we clearly cannot recover coefficients below the noise level so with more decay there will be a trade off between allowing to recover more atoms and decreasing the coefficients faster.

Theorem 2 *Assume that the signals follow the model in* (21) *and that for* $i \le S$ *the coefficients satisfy* $c_{i+t}/c_i \le 1 - \frac{\lambda}{S}$, $t, \lambda > 0$. *Then OMP will recover an atom from the support in the first s steps, except with probability* $4sK^{1-2m}$, *as long as,*

$$\left(t\left\lceil \frac{S}{\lambda} \right\rceil + \sqrt{\frac{mtS\log K}{\lambda}} + 1 \right)\left(\mu + 2s\mu^2 \right) \le \frac{1}{10} \tag{22}$$

$$and \qquad c_s \ge 14\rho\sqrt{m\log K}. \tag{23}$$

Proof We use the same approach as before, assuming w.l.o.g. $I = \mathbb{S}$, but take into account the new expression for the residuals

$$\tilde{r}_J = Q(\Phi_J)\tilde{y} = Q(\Phi_J)(y + \eta) = r_J + Q(\Phi_J)\eta \tag{24}$$

and inner products $\langle \varphi_k, \tilde{r}_J \rangle = \langle \varphi_k, r_J \rangle + \langle \varphi_k, Q(\Phi_J)\eta \rangle$. If J is an admissible sub-support, $J = A_i \cup B$ for $B \subseteq M_i$, then in analogy to (11) a sufficient condition for OMP to pick another admissible support in the noisy case is that

$$|\langle \varphi_i, \tilde{r}_J \rangle| > \max_{k \in Z_i} |\langle \varphi_k, \tilde{r}_J \rangle|. \tag{25}$$

Since we have $|\langle \varphi_i, \tilde{r}_J \rangle| \ge |\langle \varphi_i, r_J \rangle| - |\langle \varphi_i, Q(\Phi_J)\eta \rangle|$ as well as $|\langle \varphi_k, \tilde{r}_J \rangle| \le |\langle \varphi_k, r_J \rangle| + |\langle \varphi_k, Q(\Phi_J)\eta \rangle|$ for all $k \in Z_i$ the condition in (25) is implied by having for all $k \in Z_i$

$$|\langle \varphi_i, r_J \rangle| - |\langle \varphi_k, r_J \rangle| > |\langle \varphi_i, Q(\Phi_J)\eta \rangle| + |\langle \varphi_k, Q(\Phi_J)\eta \rangle|. \tag{26}$$

7

This means that we need to bound $|\langle \varphi_k, Q(\Phi_J)\eta \rangle|$ for all $k \notin J$. Using the decomposition $J = A_i \cup B$ we can rewrite

$$
\begin{aligned}
\langle \varphi_k, Q(\Phi_J)\eta \rangle &= \langle \varphi_k, Q(\Phi_J)[P(\Phi_{A_i}) + Q(\Phi_{A_i})]\eta \rangle \\
&= \langle \varphi_k, [\mathbb{I}_d - P(\Phi_J)]Q(\Phi_{A_i})\eta \rangle \\
&= \langle \varphi_k, Q(\Phi_{A_i})\eta \rangle - \langle \varphi_k, P(\Phi_J)Q(\Phi_{A_i})\eta \rangle \\
&= \langle \varphi_k, Q(\Phi_{A_i})\eta \rangle - \langle (\Phi_J^\star \Phi_J)^{-1}\Phi_J^\star \varphi_k, \Phi_J^\star Q(\Phi_{A_i})\eta \rangle. \quad (27)
\end{aligned}
$$

Since $\langle \varphi_j, Q(\Phi_{A_i})\eta \rangle = 0$ for all $j \in A_i$ this yields for all $k \notin J$

$$
\begin{aligned}
|\langle \varphi_k, Q(\Phi_J)\eta \rangle| &\le |\langle \varphi_k, Q(\Phi_{A_i})\eta \rangle| + \|(\Phi_J^\star \Phi_J)^{-1}\|_{2,2} \cdot \|\Phi_J^\star \varphi_k\|_2 \cdot \|\Phi_B^\star Q(\Phi_{A_i})\eta\|_2 \\
&\le |\langle \varphi_k, Q(\Phi_{A_i})\eta \rangle| + \frac{\mu\sqrt{|J|}}{1 - \delta_J} \cdot \sqrt{|B|} \cdot \max_{j \in B} |\langle \varphi_j, Q(\Phi_{A_i})\eta \rangle| \\
&\le \max_{k \notin A_i} |\langle Q(\Phi_{A_i})\varphi_k, \eta \rangle| \left(1 + \frac{\mu\sqrt{|J| \cdot |B|}}{1 - \delta_J} \right). \quad (28)
\end{aligned}
$$

To bound $|\langle Q(\Phi_{A_i})\varphi_k, \eta \rangle|$ for all $k \notin A_i$ with high probability, we use the sub-Gaussian property of η for the marginals $\langle v_{ik}, \eta \rangle$, with $v_{ik} = Q(\Phi_{A_i})\varphi_k$. For $s \le S$ this leads to

$$
\mathbb{P}\big(\exists i \le s, k : |\langle Q(\Phi_{A_i})\varphi_k, \eta \rangle| \ge \theta_\eta \big) \le 2sK \exp\left(-\frac{\theta_\eta^2}{2\rho^2} \right),
$$

where we have used that $\|v_{ik}\|_2 \le 1$. We substitute this bound into (28) to get, that except with probability $2sKe^{-\theta_\eta^2/(2\rho^2)}$ for all $k \notin J$

$$
|\langle \varphi_k, Q(\Phi_J)\eta \rangle| \le \theta_\eta \left(1 + \frac{\mu\sqrt{|J| \cdot |B|}}{1 - \delta_J} \right). \quad (29)
$$

Using the intermediate bound in the noiseless case from (20) as well as the estimate above with $|B| \le t\lceil \frac{S}{\lambda} \rceil$ and $\theta_\eta = 2\rho\sqrt{m \log K}$ we get, except with probability $4sK^{1-2m}$, for all $J = A_i \cup B$ with $|J| < s$ and all $k \in Z_i$ that

$$
\begin{aligned}
c_i^{-1}|\langle \varphi_i, \tilde{r}_J \rangle| - |\langle \varphi_k, \tilde{r}_J \rangle| > 0.63 - 2 &\left(t\left\lceil \frac{S}{\lambda} \right\rceil + \sqrt{\frac{mtS \log K}{\lambda}} + 1 \right) \cdot (\mu + 2\mu^2 s) \\
&- \frac{4\rho}{c_s}\sqrt{m \log K} \left(1 + 2\mu\sqrt{st\left\lceil \frac{S}{\lambda} \right\rceil} \right). \quad (30)
\end{aligned}
$$

Finally, observe that the condition in (22) guarantees that $4\mu^2 st\lceil \frac{S}{\lambda} \rceil \le \frac{1}{5}$. Thus, in order for the right hand side above to be larger than zero, which implies recovery of a correct atom in the first s steps, it suffices that

$$
\frac{4\rho}{c_s}\sqrt{m \log K} \left(1 + \sqrt{\frac{1}{5}} \right) \le 0.43 \quad (31)
$$

which is guaranteed by (23). ∎

Figure 1: Percentage of correctly recovered supports for noiseless signals with various sparsity and coefficient decay parameters via BP (a,d), OMP (b,e) and thresholding (c,f) in the Dirac-DCT dictionary (a,b,c) and the Dirac-DCT-random dictionary (d,e,f).

4. Numerical Simulations

To see how well our results predict the performance of OMP, we conduct recovery experiments both with noisy and noiseless signals in $\mathbb{R}^d$ for $d = 128$. The signals follow the model in (3) resp. (21). The permutation p is chosen uniformly at random and the sparse coefficients form a geometric series with parameter α, meaning $c_i = \beta_S \alpha^i$ for $i \leq S$ and zero else, with β_S a constant ensuring $\|c\|_2 = 1$. We vary α between 0.75 and 1 and the sparsity level S between 2 and 48. In case of noise, we choose η_i i.i.d Gaussian with variance $\rho^2 = \frac{1}{256d}$ and $\rho^2 = \frac{1}{16d}$, corresponding to signal to noise ratios (SNR) of 256 and 16. As dictionaries we use the union of the Dirac and DCT bases (coherence $\mu = 0.125$) and the Dirac-DCT dictionary with additional $2d$ vectors choosing uniformly at random from the sphere ($\mu = 0.366$).

In the first experiment we draw $N = 1000$ permutations p and sign sequences σ and for each pair (S, α) count how often BP, OMP and thresholding can recover the full support from the corresponding signals. From the results in Figure 1 we can see that the success region of OMP is indeed a union of two areas, one derived from the worst case analysis, $S\mu \lesssim 1$, and one from the average case analysis $S\mu^2 \log K \lesssim 1 - \alpha$. In particular, we can see the linear dependence of the breakdown sparsity level on the parameter α. We can also see that the success of BP is not influenced by the coefficient decay and that, as indicated by theory, neither BP nor OMP is better in general but that each of them is better in a certain region. Finally, observe that the price you have to pay for the computational lightness of

Figure 2: Percentage of correctly recovered atoms before recovery of first wrong atom via OMP for signals with various sparsity levels and coefficient decay parameter contaminated with no noise (a,d) or Gaussian noise corresponding to SNR = 256 (b,e) and SNR = 16 (c,f) in the Dirac-DCT dictionary (a,b,c) and the Dirac-DCT-random dictionary (d,e,f), as well as the percentage of correctly recoverable atoms for SNR = 256 and SNR = 16 (g,h).

thresholding is the very limited range of parameters, where it is performing well.

In the second experiment we additionally draw N noise-vectors to create the signals. For each signal we count how many atoms OMP identifies correctly before recovering the first incorrect atom. Figure 2 shows the average over all N realisations divided by the correct sparsity level for both dictionaries and three noise levels, as well as the relative number of recoverable atoms for the two non-zero noise levels (g, h), meaning the number of coefficients above the noise level $c_i^2 \geq 2\rho^2 \log K$. Comparing to the success rates in the noiseless case, we can clearly see the overlay of the two effects; for small coefficient decay we recover as many atoms as in the noiseless case, while for large decay we recover all atoms with coefficients above the noise level.

5. Discussion

We have shown that OMP is successful with high probability if the coefficients exhibit decay and in such settings can even outperform BP. In particular, for geometric sequences with parameter $\alpha < 1$ the admissible sparsity level scales as $S\mu^2 \log K \lesssim (1 - \alpha)$. Our next goal is to extend the results to OMP using a perturbed dictionary, which is a necessary step to help tackle dictionary learning algorithms like K-SVD theoretically. We are also interested in deriving average case results for other algorithms such as stagewise OMP, [5], which picks more than one atom in each round, or Hard Thresholding Pursuit, [6], an iterative thresholding scheme. Both these algorithms can be computationally more efficient due to using less iterations and a theoretical analysis might allow the design of hybrids that automatically adapt to the decay, retain computational efficiency and allow for multiple stopping criteria.

Acknowledgments

References

[1] E. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006.

[2] Emmanuel J. Candès and Terence Tao. Near optimal signal recovery from random projections: universal encoding strategies? *IEEE Trans. Inform. Theory*, 52(12): 5406–5425, 2006.

[3] D. Donoho and M. Elad. Optimally sparse representation in general (non-orthogonal) dictionaries via ℓ_1 minimization. *Proc. Nat. Aca. Sci.,*, 100(5):2197–2202, March 2003.

[4] D.L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4): 1289–1306, 2006.

[5] D.L. Donoho, Y. Tsaig, I. Drori, and J.L. Starck. Sparse solution of underdetermined systems of linear equations by stagewise orthogonal matching pursuit. *IEEE Transactions on Information Theory*, 58(2):1094–1121, 2012.

[6] S. Foucart. Hard thresholding pursuit: An algorithm for compressive sensing. *SIAM Journal on Numerical Analysis*, 49(6):2543–2563, 2011.

---

3. For $J = A_i \cup B$ in general $P(\Phi_J)Q(\Phi_{A_i}) \neq P(\Phi_B)Q(\Phi_{A_i})$.

[7] J. J. Fuchs. Extension of the pisarenko method to sparse linear arrays. *IEEE Transactions on Signal Processing*, 45(2413-2421), October 1997.

[8] A.C. Gilbert and J.A. Tropp. Signal recovery from random measurements via Orthogonal Matching Pursuit. *IEEE Transactions on Information Theory*, 53(12):4655–4666, December 2007.

[9] R. Gribonval, H. Rauhut, K. Schnass, and P. Vandergheynst. Atoms of all channels, unite! Average case analysis of multi-channel sparse recovery using greedy algorithms. *Journal of Fourier Analysis and Applications*, 14(5):655–687, 2008.

[10] C. Herzet, A. Drémeau, and C. Soussen. Relaxed recovery conditions for OMP/OLS by exploiting both coherence and decay. *IEEE Transactions on Information Theory*, 62(1):459 – 470, 2016.

[11] D. Hsu, S.M. Kakade, and T. Zhang. A tail inequality for quadratic forms of subgaussian random vectors. *Electronic Communications in Probability (arXiv:1110.2842)*, 17 (14), 2012.

[12] D. Needell and R. Vershynin. Uniform Uncertainty Principle and signal recovery via Regularized Orthogonal Matching Pursuit. *Foundations of Computational Mathematics*, DOI: 10.1007/s10208-008-9031-3.

[13] Y. Pati, R. Rezaiifar, and P. Krishnaprasad. Orthogonal Matching Pursuit: recursive function approximation with application to wavelet decomposition. In *Asilomar Conf. on Signals Systems and Comput.*, 1993.

[14] K. Schnass and P. Vandergheynst. Average performance analysis for thresholding. *IEEE Signal Processing Letters*, 14(11):828–831, 2007.

[15] J.A. Tropp. Greed is good: Algorithmic results for sparse approximation. *IEEE Transactions on Information Theory*, 50(10):2231–2242, October 2004.

[16] J.A. Tropp. On the conditioning of random subdictionaries. *Applied and Computational Harmonic Analysis*, 25(1-24), 2008.